BMVA News

The Newsletter of the British Machine Vision Association and Society for Pattern Recognition

Volume 8 Number 5
May 1998

Editor: Dr Paul Rosin
Department of Information Systems & Computing
Brunel University
Middlesex UB8 3PH
Tel: (01895) 203392
Fax: (01895) 251686
Email: Paul.Rosin@brunel.ac.uk

BMVA News[1] is published every three months. Contributions on any activity related to machine vision or pattern recognition are eagerly sought. These could include reports on technical activities such as conferences, workshops or other meetings. Items of timely or topical interest are also particularly welcome; these might include details of funding initiatives, programmatic reports from ongoing projects and standards activities. Items for the next edition should reach the editor by 17th April 1998.

Contents

[1]The British Machine Vision Association and Society for Pattern Recognition is a Company limited by guarantee, No. 2543446, registered in England and Wales. Registered Office: Granta Lodge, 71 Graham Road, Malvern, WR14 2JS. The Association is a non-profit-making body and is registered as charity No. 1002307. See http://www.bmva.ac.uk/ for more information about the BMVA.

Editorial

A few issues back I was complaining about the lack of humour in our community. Recently, someone sent me some maths jokes, which prompted me to make a quick test to compare the fields by running a Web search. Whereas I found literally hundreds of relevant matches to "maths humor" there was not a single one on "computer vision humor" (I'd given up on UK humour, and was hoping for better things abroad, hence the spelling!). In fact, emphasising the level of interest, there were websites dedicated to providing pointers to other maths humour sites. So why is that apparently mathematicians enjoy a joke but we don't?

Anyway, on the subject of humour, it got me thinking about its application to educational purposes in computer vision. For instance, if you want to learn about management, you can be guided by Winnie-the-Pooh, in the appropriately titled "Winnie-the-Pooh on Management", by Roger Allen. Other examples abound, but I'll give just a few instances of humour or whimsey from the "Mathematical Universe" by William Dunham which I rather liked. In explaining number theory he says

> the natural numbers resemble fish out of water. Spawned by the process of addition, they find themselves in an unfamiliar, multiplicative environment.

And then, on the congruency of base angles of isosceles triangle he says

> Thales' [...] actual proofs disappeared long ago. Still the ancients held him in very high regard, classifying him as one of the "seven wise men of antiquity".

(There is no truth to the rumour that the other six were Grumpy, Happy, Dopey, Sneezy, Doc, and Bashful.)

On relating these to a colleague he remarked that it is merely indicative of my infantile humour! Maybe so, but I'm still waiting for a computer vision text book along these lines. Attention prospective authors, I'm sure there's an unexploited gap in the market there.

Industrial Exploitation

As you are probably aware, the British Machine Vision Association, while primarily having stronger links with the academic community, has always had a strong commitment to facilitate vision technology transfer. In recent years this has included a leading role in the developments of funding initiatives such as the EPSRC IMV. Also, last year we introduced an Industrial Day to the BMVC in order to try to attract industrialists who otherwise may not want to attend the entire conference. Such opportunities for the academic and research communities to meet will doubtless become more important in the future, as funding for research projects becomes ever more tightly bound to exploitation.

Recently, the BMVA committee have been busy constructing a new web site advertising our activities and providing access to resources for teaching and research. This site will soon be publically available and we would like to include pages or links to pages which demonstrate the use of Machine Vision in industry. We are therefore interested in any material which people may have which can be used for this purpose. We appreciate that in many cases industrialists do not wish to explain their technological advantages to their competitors, on the other hand there may be some companies out there who would welcome an opportunity for what would effectively be free publicity. Alternatively there may be some consultants who, no longer bound by confidentiality agreements, are in a position to demonstrate their talents by providing examples of successful technology transfer in a real application. In any of these cases we would like to hear from you.

Please send me an email at nat@sv1.smb.man.ac.uk with any suggestions and we will get back to you.

Neil Thacker
University of Manchester
Department Medical Biophysics
email: nat@sv1.smb.man.ac.uk

Call for Speakers

BMVA Technical Meeting

Visual Guidance of Robots

Wednesday, 15 July 1998

In recent years, vision algorithms in areas such as object recognition and 3D reconstruction have matured to the point where it is now a practical proposition to revisit one of the original motivations for vision research – the use of visual information to autonomously guide robots.

The intended purpose of this one-day meeting is to examine state-of-the- art approaches to the autonomous visual guidance of robots, where a robot is broadly interpreted as any mechanical contrivance: robot arm, land or sea or air vehicle, or manufacturing device of significant complexity. Contributions of original work (and work in progress) are invited which cover the application of vision as a sensory input for robotic devices to produce autonomous action. Contributions which combine vision with other measurements – especially uncertain measurements such as odometry – are particularly welcome.

Additionally, much interesting work is currently being undertaken in virtual reality and germane contributions are also invited which address the construction of virtual worlds from real worlds and link to the concept of an action.

This meeting is jointly sponsored by BRA (formerly known as the British Robot Association) and it is hoped this meeting will bring together academic researchers working in vision and related fields, with developers and end-users of robots.

Readers should contact Peter Rockett IMMEDIATELY if they are interested in speaking at this meeting:

Dr Peter Rockett
Dept of Electronic & Electrical Engin.
University of Sheffield
e-mail: p.rockett@shef.ac.uk

Vision in the Built Environment

The BMVA technical meeting was held on Wednesday the 19th of February chaired by Tim

Ellis of City University and Gordon Mair of Strathclyde. The focus of the meeting was the use of computer vision techniques for visualisation and reconstruction of built environments with particular emphasis on the construction industry.

Dave Chapman from UCL described a project involving the creation, management and exploitation of massive image databases of process plant facilities. The main application is reverse engineering to create "as built" CAD models of the plant for planning new building work. A motorised camera head creates highly detailed panoramic images which are then used to extract CAD data using photogrammetry techniques. The main problems are the huge amount of data, the complexity of the environment and the model conversion which is time consuming and essentially manual.

Tony Pridmore from Nottingham University described a technique for the interpretation of video image sequences inside small bore pipes. An optic flow approach has been adopted but the classical algorithm of Horn was found to be unsuitable as the vector field smoothness and fixed lighting assumptions are both violated. The smoothness constraint was changed to admit vectors of similar direction rather than magnitude giving much more robust results. The pipe is lit with a moving light source which effectively creates a second flow field multiplied together with the texture flow field. The fields are separated using a simple but ingenious method. A log transform of the input is taken in which the fields are additively related. Gaussian smoothing is performed and the resultant image, largely the lighting component, is subtracted from the input.

Christine Baillard described progress at Oxford University on the automatic extraction of buildings from aerial photographs. Robust automatic extraction is difficult as the scene detail is complex, varied and dense. They have found that great improvements can be made by integrating information across more than the usual two views using the trifocal tensor as a unifying framework. An improvement from 48% to 68% of edge segments correctly matched was recorded when integrating up to six images. Future work will further automate the process of 3D model building which is currently semi-automatic.

After lunch, Kia Ng of Leeds University described results from the ACTS RESOLV project. A dense laser depth scan is combined with video images to provide a textured VRML model of an interior for telepresence applications. This system can be mounted on a mobile robot which can autonomously plan successive optimal positions for data capture. A single scene can now be captured in 10 minutes as opposed to 30 minutes previously. Demonstrations were also given of registering movies within the virtual scene.

Gordon Mair from Strathclyde University described work on a system for telepresence in the building industry. The system combines components from telepresence, VR and telecommunications. A stereo head is slaved to motion of a VR headset where control signals and compressed video are sent via mobile phone. Two phones are used to provide a total bandwidth of 18K bytes per second. Although the links are relatively low bandwidth the system can potentially be used in any remote site and impending improvements in mobile telecommunications will improve the performance of the system.

After a tea break there were two presentations from representatives from industry.

Denis Chamberlain, a civil engineer from City University, described the role of computer vision in construction and repair. He said there have been few success stories so far for reasons within both the research and construction communities. Researchers must remember the harsh and unstructured environment when designing systems which have to aquire safety permits for use on site. It is also best to create robust stand alone products. There is little interest in hardware development within the industry where most emphasis is placed on using IT for direct cost savings. Interesting trends include offsite prefabrication (more structure for vision systems), smart intelligent buildings (contact architects) and mandatory health and safety legislation spurring safety monitoring and CDM (construction design management). One example of a successful vision application outlined was a system for monitoring, planning and coordinating crane movements.

The day was rounded off by a talk by Nicholas Farrow from BICC. The company is involved in many initiatives including CICC (Collaborative Integrated Communication for Construction). The objective of CICC is to allow collaboration of dispersed and temporary teams. An important goal is cost savings through avoiding mistakes which occur through poor communication between contractors. An augmented reality approach is adopted which combines the realism of imagebased VR with the flexability of model based VR. Impressive demos from the Fraunhofer Institute of registering a virtual bridge designs with images and video were shown. Again discrete tools which work well were preferred to complex integrated solutions as this approach does not constrain implementation and delivery.

Co-chair Tim Ellis wrapped up what had been an

interesting meeting. Clearly, computer vision is set to play an expanding role within the construction industry and in the reconstruction and visualisation of built environments.

Mark Wright
Edinburgh Virtual Environment Center
University of Edinburgh
email: Mark.Wright@ed.ac.uk

Image and Video Databases

The BMVA's technical meeting last December was entitled "Image and Video Databases" and brought together over 40 researchers from both academic and industrial backgrounds, with several travelling from as far afield as France and Spain specifically to attend. Eight interesting talks were delivered and although there was a bias towards static images rather than dynamic ones, the subjects presented were in most cases relevant to both, resulting in stimulating discussion.

The day's events were kicked-off with an excellent talk from Dr Colin Bird of IBM's UK Laboratories in Hursley Park. He discussed the potential applications to which the fruits of purer research may be put, using his experience with IBM's QBIC system as his focus. Colin raised two interesting points regarding searching by content: (a) Can it actually be done? (b) Will we actually find it useful when if it is possible! The first point is valid when one considers some of the complex feature sets being used to represent query keys – how will an ordinary user, with no technical background be able to control such a system? It was suggested that perhaps we should step back and think about how our systems would be employed. Is one looking for a specific image in a large collection, such as a home photograph database? In this case we should focus on relevance. Alternatively, one could be seeking creative inspiration, perhaps for textiles design, where similarity is more important.

Moving from high-level applications, to low-level features, the next speaker was Sadegh Abbasi from Surrey University, who talked about a representation for the shapes of segmented objects using Curvature Scale Space. This method starts with the outline of the object and progressively smoothes it, plotting out the curve's maxima, resulting in a series of 'loops', the characteristics of which can be plotted in 2D space. This forms the basis for a database search by comparison of these plotted points. Some results were demonstrated and were quite impressive despite being currently restricted to binary images. One could imagine how this technique could be combined with colour, texture and other features in a less constrained domain to create a powerful search tool.

Stephanie Fountain from Reading University discussed a database annotation and retrieval tool known as RAIDER which she is designing for her PhD. Of particular note in this talk was the technique of user-guided segmentation in which a bounding box is employed to designate the bulk of a region to be segmented. The system (hopefully) fills in the remainder, using texture analysis as its metric. Objects thus segmented are then labelled in a combination of manual and automatic tagging. Those which are manually tagged are used as keys to find other similar objects in the database so that the same label can be applied to them, with the user's confirmation. The hope is that the more the system learns about objects, the less user interaction will be required to tag new images until the process becomes fully automatic.

The next talk was given by Benoit Huet who discussed the use of shape for retrieval of line-patterns. The database used consisted of aerial photographs and company logos segmented to highlight the linear features. Two and three dimensional shape histograms were obtained and various properties used for comparison, such as pairwise angular difference, length ratio and cross ratios. Some results were presented and demonstrated the technique's effectiveness. The comparison of every line with every other line in the image was shown to be rather a large task and was simplified by comparing just n nearest neighbours without adversely affecting the results. Despite this reduction in complexity, the system was shown to be robust to the quality of segmentation.

After a pleasant lunch spent discussing the morning's talks, I presented some work I have been doing in conjunction with Hewlett Packard Labs. This project has been looking at content-based retrieval from the relatively unconstrained domain of holiday photographs. My approach employs user-feedback to iteratively refine the query results, rather than attempting to perform a one-shot query with complex feature sets. I demonstrated some results which (I hope!) showed how effective this method can be.

I was followed by Dr Adrian Thomas from Essex University who managed to use more slides than I have ever seen in a 25 minute talk! His background in 3D modelling has led him to consider extracting 3D shape from objects in video clips to aid recognition,

retrieval and tracking. The project is in it's early stages but some shows a lot of promise.

James Shanahan from Bristol University then gave a talk on an interesting approach to the classification of objects in segmented scenes using Fuzzy Cartesian Granule Features to cross-correlate characteristic features from regions. These correlations form fuzzy sets which can be used to semantically identify objects in constrained domain images with a high degree of accuracy.

Dr Mark Dunlop from Glasgow University provided an excellent close to the day by discussing measures which can be employed to assess the quality and ability of query system to live up to the user's expectations. He proposed that although automatic indexing techniques are much less labour-intensive and potentially more consistent, they are often less accurate than manually created indexes. Forging a clever link with Colin's talk, he showed how systems can be tuned for precision (i.e. looking a specific image), or recall (i.e. looking to retrieve as many images of a given type) despite the fact that the user may well require both. A highly poignant statistic was given based on a survey of client requests to photograph library agencies. This showed that features which are traditionally used for indexing in automatic systems rarely appeared in the client's descriptions with much more subjective and contextual information forming the bulk or requests. Mark suggested that in order to usefully employ this form of meta-data, it will have to be separately produced and attached in some manner to the image which is hopefully where MPEG-7 will help. A spontaneous discussion followed Mark's talk and it was pleasing to see the unusual situation of the majority of delegates agreeing with each other! The general feeling is that automatic digital image and video retrieval is very much in its infancy and that there is still a long way to go. However, there were some overtones of caution which hinted that we don't fully understand what an ordinary user may wish to actually do with such systems and indeed whether they will ever be able to completely fulfill their aims. Is a bright green steam train (or was it a red one?) just too difficult to represent without manually created meta-data? Overall, the day was very successful and enjoyable and my thanks go to all the speakers for giving their time and to all the delegates for attending and contributing to the discussions.

Matthew Wood
ACRC Vision Group
University of Bristol
email: `Mat.Wood@bris.ac.uk`

Book Review

Motion-Based Recognition

Edited by Mubarak Shah and Ramesh Jain (Computational Imaging and Vision series, Kluwer Academic Publishers), this text documents the state of the art (in 1996!) in motion-based recognition with a collection of invited chapters from leading researchers in this rapidly emerging field. Following an introductory survey by the editors, the book is divided into three main subject areas: human activity recognition, gesture and facial expression recognition, and lipreading.

Part I describes several approaches to the problem of detecting and recognising human activities, ranging from the low-level scale-space based approach of Yacoob and Davis for computing optical flow and acceleration, to the high-level 3D cylinder model of Rohr which uses joint curves derived from motion studies. Another low-level approach, change detection, is used by Baumberg and Hogg for automatically generating deformable contour models of the human body (based on the PDM and demonstrated by tracking in real-time), and by Bobick and Davis to generate Motion Energy Image and Motion History Image templates which are used to recognise actions. An interesting characteristic of many natural motions, such as walking or a heartbeat, is their cyclic or near-periodic nature. This characteristic is exploited by Seitz and Dyer, using an interesting view-independent, compact description of near-periodic motions, by Polana and Nelson, who consider motion from three categories based on the extent of any spatial or temporal repetition, and by Goddard who studies structured motions using a biologically inspired connectionist approach.

Part II describes four approaches to gesture recognition and facial expression recognition. Both approaches to gesture recognition are essentially HMM based. Bobick and Wilson present a HMM-like approach where gesture and its variability are described by a sequence of states generated from prototype trajectories in a configuration space, whilst Starner and Pentland use a number of small HMMs of fixed topology to model American Sign Language from relatively crude hand tracking. For facial expression recognition, Black et al. use a parameterised model of the optical flow of different facial patches, whilst in Essa and Pentland's approach, the state of a dynamic, physics-based face model is computed from optical flow, resulting in interesting motion-energy templates for a number of expressions.

Part III describes three approaches to visual speech recognition or lipreading. As with gesture recognition, two of the approaches are based on HMMs. In the approach of Bregler and Omohundro, the HMM is embedded in a learnt 'lip manifold' in shape space, whilst Goldschen et al. use an HMM in an attempt to recognise continuous speech based on oral-cavity features. Finally, Nan et al. present an Eigenvector representation of image sequences for different spoken letters.

All chapters are well written with plenty of descriptive diagrams. The book will be of particular interest to researchers in Human-Computer Interfaces and Computer Vision, although it may also be of interest to researchers in Psychology, Computer Graphics and Multimedia.

Neil Johnson
Leeds University
email: neilj@scs.leeds.ac.uk

Book Review

Human Symmetry Perception and its Computational Analysis

edited by *C. W. Tyler*

published by VSP, Utrecht, 1996
393 pages, paperback,
ISBN 90-6764-210-X

Who does not remember diving into the visual perception literature, only to shrug at the sight of the simplistic input images and the rows of plots, followed by apparently unrelated conclusions about the way human vision is really implemented?

At a first glance this collection of 22 recent papers[2] on the perception of bilateral symmetry may be subject to the same impression, but its strength lies in the focus on a single phenomenon, and the diversity of the approaches to the subject. Following a motivating introduction by the editor, the remaining papers are divided into two parts: (1) *empirical evaluation* and (2) *theoretical issues.*

The first part comprises eight papers describing psycho-physical experiments measuring specific performance characteristics of human symmetry perception. They are accompanied by an excellent review of empirical evidence and explanatory theories by Wagemans, and between them they offer more background material and coverage than one might expect from a collection of papers. In addition the papers are generally well written, and, being original research papers, they offer sufficient detail to enable an outsider to this field to understand and interpret the conclusions drawn by the authors.

The twelve papers comprising the second part are concerned with theoretical issues, mainly in the form of computational models of symmetry detection. Contrary to the work published in the computer vision community, which invariably is application driven, these papers are more concerned with the question of how well the performance of a computational model agrees with human performance. The work described here includes models devised to explain existing psycho-physical data, comparison of existing symmetry detection algorithms to human performance, as well as more general analysis of the properties of images of symmetric objects (Vetter & Poggio). The contributions are quite diverse, and there is no single attempt at reviewing all the existing research on computational models and symmetry detection algorithms. But this same diversity makes the combined background material and bibliographies an interesting resource.

This collection complements the computer vision literature on the detection and use of symmetry. It will be of interest to researchers working on those problems, as well as to anyone wanting to add to his knowledge about visual perception, and current research in that area. A collection of papers will never be as accessible as a good textbook, but until one is written, this may be as close as you can get.

Torfi Thorhallsson
University of Oxford
email: torfi@robots.ox.ac.uk

[2]Most of the papers originally appeared in one of two special issues of *Spatial Vision* (Vol. 8, No. 4, 1994 and Vol. 9, No. 1, 1995).

Book Review

Indirect Perception

by Irvin Rock,

published by Bradford, 1997,
ISBN 0-262-18177-0

'Indirect Perception' is an edited volume of experimental work and a few essays, by the late

Irvin Rock. A forward by Stephen Palmer sets the scene for much of what is to follow by remarking on Rock's views in relation to 'inferential and computational', 'gestalt', and Gibsonian, 'direct' approaches to human visual perception. In twenty three chapters, more than half of which Rock contributes himself, direct perception–the notion that perception is determined by the pattern of light falling on the retinae, with no intervening processing–is repeatedly challenged through the results of many, often deceptively simple, psychophysical experiments. In several general introductions by Rock, as well as within the chapters themselves, the logic and development of the experiments are clearly explained providing an easy and pleasurable read. One theme that emerges repeatedly is that "at least some perceptions are ... determined by other perceptions", a situation that, according to the arguments presented in the book, undermines 'direct perception'.

An alternative perspective on many of the experiments is what Rock and Palmer call the "levels question". This is to ask at what level of processing does a particular visual computation,the perception of lightness for example, occur. While many computational theories often suppose that much visual processing can take place at an early level–the level of the two-dimensional retinal image for example – experimental results repeatedly show that this level is, in general, too early, and that attribution of surface reflectance, element grouping, and other computations, must (also) happen at a later stage, after three-dimensional properties of the viewed scene have been assigned.

Other topics covered in the volume are: amodal completion; perception of shape, symmetry and motion; speed and size constancy; masking and visual illusions.

I would recommend this book to anyone interested in visual perception, particularly those who appreciate the power of simple demonstrations and arguments.

Tim Meese
Vision Sciences,
Aston University.
email: `t.s.meese@aston.ac.uk`

Book Review

Early Visual Learning

edited by Shree K. Nayar and Tomaso Poggio

published by Oxford University Press, 1996
ISBN: 0-19-509522-7

Early Visual Learning is a collection of thirteen papers. At a first glance, the individual contributions appear to have almost nothing in common besides addressing the problem of visual learning. The subjects are diverse indeed, ranging from a study of monkeys' ability to recognise 3D objects from previously unseen views to a description of a system for counting volcanos on Venus. But this edited volume is not just a pot-pourri of related articles. Neither the title nor the preface stress this, but the book has a lot to offer to researchers interested in recognition. With few exception the assembled papers describe systems that 'learn to recognise'. In particular almost half of the book is devoted to view-centered, appearance-based techniques to object recognition.

The second common feature of the contributions is not immediately apparent either. All the presented learning systems heavily rely on methods based on statistics, pattern recognition and function approximation. The reader thus gets a good overview of representations, techniques and methods well established and tested mainly in pattern recognition systems that are directly applicable in vision. 'Gaussian mixture models', 'principal component analysis', 'the EM algorithm', 'genetic algorithms' are names that spring to my mind.

Edited volumes are often criticised from the point of view that most of the work presented in the volume has been published previously. This is true for 'Early visual learning' too, but the contributions are not re-printed version of conference papers. This can be illustrated by Moghaddam's and Pentland's contribution describing the MIT face recognition system. I work in the area and I can give five references to conference and journal publications that cover the 30-odd page chapter. But by integrating, cutting out overlapping parts and polishing the material the authors have added enough value to easily justify a 'repeated' publication. In fact, together the three chapters of 'Early Visual Learning' that address the face recognition problem might be one of the best starting point for anyone who wants to get acquaint with appearance based face recognition.

Chapter 1, written by Nayar and Poggio, servers as an introduction and motivation. The authors argue that learning is a central aspect of intelligence, proposing a modified Turing test. Turning their attention form general AI system to vision, the authors discuss the role of learning in vision systems. The remaining chapters are individual case studies

of learning system.

In Chapter 2, Pauls, Bricolo and Logothetis describe combined psychophysical and electro-physiological experiments on rhesus monkeys whose object recognition performance is tested under changing viewing condition. The experiments show that recognition rate is not view-point invariant, supporting the hypothesis of a view-based representation of objects in the brain. Next, Poggio and Beymer formulate the learning-from-example problem as one of multivariate function approximation. The technique is applied to face recognition. The subsequent chapter by Pope and Lowe describe a probabilistic framework for object recognition. Unlike many other contributions in the collection, recognition is based on geometric features like lines, parallel segments etc.

The already mentioned work of Moghaddam (Chapter 5) uses pattern recognition techniques such as principal component analysis and mixture of Gaussian models for Face detection, coding and recognition and gesture recognition. Nayar, Murase and Nene (Chapter 6) present a parametric eigenspace model. Object appearance is represented by a low-dimensional manifold. Recognition is achieved by closest point search, an interesting problem in its own right, for which the authors propose and efficient solution. The next chapter by Pomerleau describes the application of neural networks to the problem of vision-based autonomous navigation. The next contribution by Weng again addresses the face recognition problem. A whole range of pattern recognition tools are brought together to tackle the problem - dimensionality reduction, hierarchical feature grouping, space partitioning trees.

Chapter 9, by R. Nelson, explores the complexity of visual learning arguing in favour of indexed memory. In the following chapter Fayyad, Smyth, Burl and Perona address the issue of learning in the context of very large image database (in the order of terabytes), where efficient data processing, classification and data management is essential. The applications are the techniques are in the area of astronomy, e.g. the volcanos-on-Venus problem mentioned above. Chapter 11 by Bhanu, Wu and Lee is concerned with image segmentation. In this case, learning is introduced to replace manual fine-tuning of parameters by automatic adaptation that uses a combination of genetic algorithms and hill climbing. The penultimate contribution by Greenspan returns to the recognition problem, in the form of learning and classifying textures. The final chapter by Salganicoff, Rucci and Bajcsy departs from mainstream computer vision addressing the problem of learning simultaneously from visual and tactile information in the context of robot control.

To conclude I can say that I enjoyed reading most of the chapters; others were too far from my interests. In such a diverse collection this is natural; I doubt many will read the book cover to cover. But for readers interested in learning in vision, (appearance-based) object recognition, application of pattern recognition in computer vision and related fields of study the book may be a valuable source of information and inspiration And, unlike many edited volumes, the index is there to guide you.

Dr. George Matas
CVUT Praha
Centre for Machine Perception
Czech Republic
email: matas@cmp.felk.cvut.cz